

The NEXXT Perspective

In an industry built on spectacle, we're often guilty of chasing the "because we can" moment – the demo that pushes technology right to the edge of what's possible. But here's the truth we don't talk about enough: capability alone doesn't earn trust. And trust is what makes technology stick.

Bill Fons found this out the hard way. His greeting avatar – a pitch-perfect replica of himself – was meant to be a showcase of hospitality and technical prowess rolled into one. Instead, it triggered the uncanny valley: that flicker of hesitation, the instinctive recoil when something is almost human but not quite right.

The takeaway isn't simply "avoid too much realism." It's that audiences – whether customers, colleagues, or the general public – respond far more positively to AI when it brings a value that is *different* from a human's, not a mimicry of it. Right now, the fastest way to erode trust is to position AI as a one-for-one human replacement. We'll embrace an AI that offers speed no human could match, memory that never fades, fluency in dozens of languages, instant recall of complex histories, or the ability to process data in ways we can't. But ask us to accept it as "just like" a human, and most people will resist. Not because the technology is bad, but because the premise feels wrong.

This is the real opportunity in persona design. AI doesn't need to fool anyone into thinking it's human. The smarter play – strategically and emotionally – is the opposite: make it unmistakably AI, and give it an identity that makes its value self-evident. A persona with its own role, tone, and personality. One that complements, rather than competes with, its human counterparts.

When we do that, the conversation shifts. A support agent becomes a 24/7 expert who never gets flustered. A tour guide becomes a living archive who remembers every guest's visit in perfect detail. A service concierge becomes multilingual and instantly adaptive to different cultural contexts. These are not roles humans can't do at all, but they are roles AI can often do faster, more consistently, or at a scale no human could.

And here's the kicker: thanks to today's AI toolkits, creating those distinctive, high-trust personas no longer requires the budget or infrastructure of a flagship experience centre. Visual presence, natural voice, expressive intonation, even stylised facial or body language cues are now available in off-the-shelf tools for a fraction of what they cost just two years ago. Which means any integrator, brand, or organisation with the curiosity to experiment can start building them today.

The risk, as Bill's experience shows, isn't in failing to reach human perfection – it's in chasing it at all. The future of AI in our spaces won't belong to the most convincing clones. It will belong to the most capable companions: those designed not to replace us, but to make the moments that matter unforgettably better.



From Greeting to Gut Check: A Personal Experiment in AI and UX

We've all heard the AI promises: personalization at scale, smarter engagement, "wow factor" moments. So when I set out to create a virtual welcome greeter – an ultra-realistic version of myself to kick off our company's experience tour – I wasn't thinking about philosophical dilemmas. I was thinking about impact.

The tech worked. The face looked like mine. The voice sounded like mine. The message was clear, polished, intentional.

And yet... something felt off.

What was meant to be warm and personal landed, for some, as unsettling. Far too lifelike. And far too synthetic. That's when I stumbled across a concept I hadn't fully considered: the **"Uncanny Valley"**.

This wasn't a failure of the technology. It was a blind spot in the experience design.

And it taught me a lesson I now carry into every AI initiative: *Just because you can doesn't mean you should*. The most important part of AI innovation isn't capability – it's how it lands on the human side.

Background

In late 2023, I was the President of PTG, a well-regarded AV integrator specializing in the restaurant and retail space. At the time, we made a strategic decision to build our very first experience center – complete with a guided tour – not as a high-budget showcase, but as a testbed for what was possible using the latest low-cost AI tools and automation platforms. Our goal was to demonstrate to customers that innovation wasn't just something we talked about – it was something we lived, built into our own operations, and could replicate affordably for others.

Our experience tour was designed to immerse visitors in what AV and digital transformation could feel like when it was done right. First impressions mattered.

So I thought: *Why not greet guests with an avatar of myself?*

We had the tools. Deepfake video tech. Voice cloning. A proximity sensor. AI integrations with my calendar. We even built an automation that would:

- Detect when a guest entered the lobby
- Pull the meeting info from my schedule
- Cross-reference the guest's travel city
- Retrieve the local weather in both cities



- Compose a short custom greeting that blended hospitality with personalization

It was smart, slick, and frankly, impressive. To most visitors, it was clearly next-level.

Until it wasn't.

The Moment It Shifted

Some guests smiled. Some laughed. A few leaned in, fascinated by how "real" it looked and sounded.

But then came the side glances. The small hesitation before speaking. The comment from one guest I still remember clearly:

"That avatar was a little too real... honestly, it creeped me out a bit."

What struck me wasn't just the feedback – it was how *unexpected* it was. I had obsessed over functionality, wow factor, and automation. But I hadn't once paused to ask: *How does this make someone feel?*

Uncovering the Uncanny

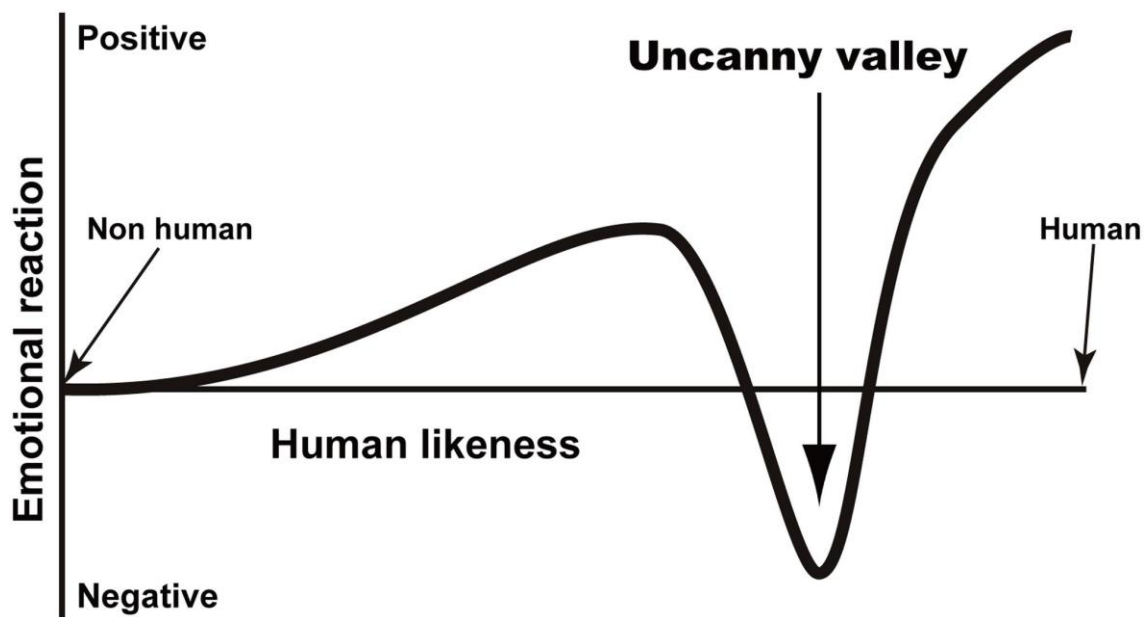
That night, my Director of Marketing, Chris Ladaucer, pulled me aside. He'd had a hunch something about the avatar wasn't landing. He did some digging and introduced me to the research around the uncanny valley.

The uncanny valley, for those unfamiliar, is a well-documented psychological and perceptual phenomenon where human-like replicas that are *almost* real – but not quite – elicit discomfort, unease, or even distrust. It's that strange gut feeling you get when something looks nearly human, but just off enough to be unsettling.

The term was first coined in 1970 by **Masahiro Mori**, a Japanese roboticist, who observed that as robots become more human-like in appearance, people's emotional response to them becomes more positive – but only up to a certain point. Once a robot or avatar becomes *too* human without being *fully* convincing, emotional affinity plummets into a "valley" of discomfort. That's the uncanny valley.

Mori illustrated this effect with a graph – familiar now to many in AI and robotics – where the y-axis measures familiarity or comfort, and the x-axis tracks human likeness. As realism increases, so does affinity – until it drops steeply. The curve only rises again once a design crosses the threshold into truly indistinguishable realism. But that threshold is steep and slippery, and most creations fall just short.





Over the years, this effect has been studied in fields ranging from robotics to animation to digital avatars. Researchers have tested facial expressions, eye movements, and speech patterns – all trying to understand why slight imperfections cause such strong reactions. The takeaway? Our brains are finely tuned to detect "almost human" cues – and they don't forgive well.

And that's what I had unintentionally built: something *too perfect to be casual*, and *too artificial to be warm*.

What I Changed

1. Avatar Realism Dialed Back

We made the digital version of me **less identical**. It still looked like me, still used my voice, still spoke in my phrasing. But there was now zero confusion: it wasn't actually me. That simple design choice removed the subtle emotional betrayal some guests were feeling.

2. Killed the Weather Script

The weather comparison – while clever – felt hollow. Automated empathy isn't empathy. I learned that even small expressions of care need to be authentic. We removed the weather lookup entirely.

In its place? A light-touch prompt that asked *me* a few hours before the meeting what I wanted the avatar to say. A simple question: "*What's a good way to welcome your guest today?*" That gave me the final word and reinserted humanity into the process.



The Bigger Lesson: Never Automate Genuine Sentiment

From "Because We Can" to "Why We Should"

The lesson I took away from this experience applies far beyond one greeting bot. It speaks to a challenge the AV industry has unintentionally created for itself: we often lead with the technology – because we can – before we've truly clarified why we should. Our solutions too often celebrate feature-rich integration and clever automation without always grounding those capabilities in real human or business value.

This pattern is especially dangerous in AI, where technical prowess can dazzle, but also alienate if it isn't solving a problem the client actually cares about. As integrators and innovators, we have to start with a clear purpose: what customer challenge are we solving, what emotion are we trying to evoke, or what trust are we trying to build?

My own misstep – automating a welcome that seemed human but lacked authenticity – reminded me that cleverness isn't connection. If our solutions don't serve a true need or elevate a moment that matters, they risk becoming distractions rather than differentiators.

Looking back, that experiment became a core philosophy for how I and others approach our AI implementations at my current company, CTI, where I serve as President of AI and Innovation.

It took some humility to admit: I had been excited by what AI could do, and blind to what it should do. We're integrators, technologists, creators. But we're also hosts, guides, and humans first.

AI can impress. But if it alienates, what's the point?

Don't use AI to imitate connection. Use it to support it.

That's a principle we should all carry as this tech becomes more embedded in our workflows.

Reflections for Integrators and Innovators

If you're building client experiences, solution tours, or even AI tools for internal use, ask yourself:

- **Are you testing for emotional impact, not just functionality?**
- **Do users clearly know when they're interacting with a machine?**
- **Does the automation support the human, or replace the part they care about?**

In our rush to impress, it's easy to forget that *experience* is an emotional journey.

Conclusion: From Clone to Companion



While I no longer work at PTG and can't say for certain if the avatar still exists, the experience left a lasting impression. For me, it served as a turning point – a reminder that tech should never outshine trust. The goal was never to create a replacement, but a thoughtful tool that supports human presence, not substitutes for it.

The uncanny valley taught me that **fidelity is not the goal. Trust is.** And you can't build trust with a synthetic smile, no matter how realistic the skin texture.

So the next time you're tempted to automate a human moment, pause and ask: *Is this about efficiency? Or is it about connection?* One of those can be faked. The other shouldn't be.

Because when it comes to presence, warmth, and welcome – the real win isn't doing more. It's showing up in the ways that matter most.

Bill Fons is a seasoned leader in the AV and AI space, and a regular contributor to the NEXXT Perspective. He currently serves as President of AI and Innovation at CTI, is a member of the AI and Cyber Security Council for the National Systems Contractor Association (NSCA), and serves on the Channel Advisory Board for Blue Touch Paper, an international AV Integration Research Organization. He brings decades of industry insight to the intersection of human experience and technology

